

MANAGED SERVICES DECLARATION TIA COMMENTS – OPEN INTERNET DOCKET

1 Introduction

My name is Marcus Weldon. My business address is 600-700 Mountain Ave., Murray Hill, NJ 07974. My current position is Corporate Chief Technology Officer of Alcatel-Lucent, and I am also a member of Bell Laboratories. In this position, I am responsible for coordinating the technical strategy across the company and driving technological and architectural innovations into the company's portfolio. I joined AT&T Bell Labs in 1995 and have won several scientific and engineering society awards for my work on electronics and optical materials. I earned a B.S. in Chemistry and Computer Science and a Ph.D. degree in Physical Chemistry from Harvard University. This declaration is prepared in support of comments filed by the Telecommunications Industry Association (TIA) with regard to the Notice of Proposed Rulemaking proposing to regulate broadband industry practices. I make the statements in this declaration based on my personal knowledge and experience as an engineer developing broadband equipment and technologies.

Alcatel-Lucent USA, headquartered in New Jersey, employs more than 17,000 people in the United States (many of them engineers), and the Company's Research and Development arm, Bell Labs, lies at the heart of the U.S. innovation economy. Alcatel-Lucent's market leadership and experience in broadband provides unique insight into policy prescriptions for the broadband era. Alcatel-Lucent has a presence in 130 countries, and has significant experience in deploying current and next generation wireless and wired broadband under a variety of geographical, regulatory, and economic conditions, for private and public entities alike.

2 Summary of Declaration

Below, I provide an overview of Managed Services: how they should be broadly defined and what needs and purposes they meet. I also describe the two current categories of Managed Service – operator-based and end user-requested – and explain how Managed Services can be provided over both IP networks and the Internet. My declaration then explains operation of Managed Services such as their traffic flow control. Finally, I discuss the potential impact broadband services regulations could have on Managed Services and the related impact on innovation and investment in the network generally.

3 Definition of Managed Services

The simplest definition of a Managed Service is a service that requires a 'better than Best Effort' delivery paradigm, with service-level guarantees.

Typical Characteristics of Managed Services

Managed Services can be defined as operator-provided or consumer-demands services that require:

- Guaranteed (low) packet loss
- Guaranteed (low) packet delay
- Secure, private connectivity
- Guaranteed bandwidth

Generically, the most common type of Managed Service is one that is highly sensitive to either packet loss or packet delay and, as such, significant impairment of the service may result if delivered with no differential forwarding relative to generic Internet or IP traffic.

Web services such as web browsing, e-mail, instant messaging (IM), social networking and on-line commerce are based on simple transactions and data transfers with no real-time service needs. For these services, if the transaction completes within even a few seconds, the user is typically satisfied and, if packets are dropped, the user often is unaware because they are resent before the end of the data transfer.

However, there are many other services delivered over the Internet or IP networks that do have real-time service delivery requirements and cannot tolerate packet loss or delay on the order of seconds, or even fractions of a second. High quality video streaming, audio and video communications and video monitoring services can be delivered as Managed Services so that they are ‘always available’ (independent of time of day and concurrent usage) and are not corrupted by generic web traffic that feature lower services requirements but which have the ability to consume the available bandwidth (e.g., large file transfers or peer-to-peer sessions).

Meanwhile, there are other types of Managed Services – services that require guaranteed security and privacy such as traditional virtual private network (VPN) services, and those that just require bandwidth guarantees, such as on-line storage or media or database back-up services. In neither of these cases are packet loss and delay necessarily of importance, although in some cases, these might also be requirements.

Common Types of Managed Offerings

Managed Services are offered in many different contexts, including consumer, enterprise, public safety, etc. Below, I list some of the most common Managed Services offered today:

- In the consumer space, the following could be delivered as Managed Services:
 - IPTV services
 - Voice and video calls
 - On-line gaming services
 - E-health (with 1 or 2 way video)
 - Remote monitoring (of the home)

- Remote access to home content
- Standard definition (SD) and high definition (HD) web video streaming with guaranteed quality
- Small office/Home Office (SOHO) communications services
- In the enterprise space the following could be delivered as Managed Services:
 - Voice and video calls and conferences
 - Inter-site SD and HD webcasts with guaranteed quality
 - Inter-site data transfers and web services with high throughput and low delay requirements (e.g., Enterprise Resource Planning (ERP) and other business-critical applications)
- In the PEG (Public, Education and Government) space the following could be delivered as Managed Services:
 - Multi-site SD and HD classroom webcasts with guaranteed quality
 - Public safety audio and video broadcasts
 - Two way audio and video ‘incident’ feeds to Police, Fire, Ambulance
 - Local community video streaming (e.g., community TV)
 - Secure, HD multimedia conferencing (for remote diagnostics, analyses, control)

Common Features to All Managed Services

In today’s market place, Managed Services fulfill the needs of quality of service, bandwidth availability, and security. As noted above, Managed Services are those services that are sensitive to packet loss or delay, or require secure ‘virtual private’ connectivity or bandwidth guarantees. This set of attributes is common to Managed Services, although not all of these attributes are required for all Managed Services. For example, audio communications services require minimal packet delay. They typically also require a minimum bandwidth guarantee, but do not have stringent packet loss requirements. In contrast, video communications typically require low variation in delay (called ‘jitter’), low packet loss and a (higher) minimal bandwidth guarantee. Meanwhile, VPNs may only have privacy/security and bandwidth guarantees, but no specific loss or delay guarantees.

The absence of such guarantees for voice or video services will result in unacceptable variation in service, such as break-up of the audio or video signal, causing clipping or even disconnection of a voice call, or pixelation or freezing of a video signal. For VPNs, the absence of such guarantees would mean that mission critical applications or communications would not always be available with the required responsiveness or with security from outside sources or third parties.

A Managed Service is Distinct from Other Internet or IP-based Offerings

For the sake of clarity, it is important to make the distinction between Internet Access and other IP-based offerings. Today, many IP-based services such as IPTV or VoIP, and some Enterprise VPNs are provided as Managed Services that do not traverse the Internet. Rather, they are delivered entirely using the network operator's IP network and associated leased lines or facilities.

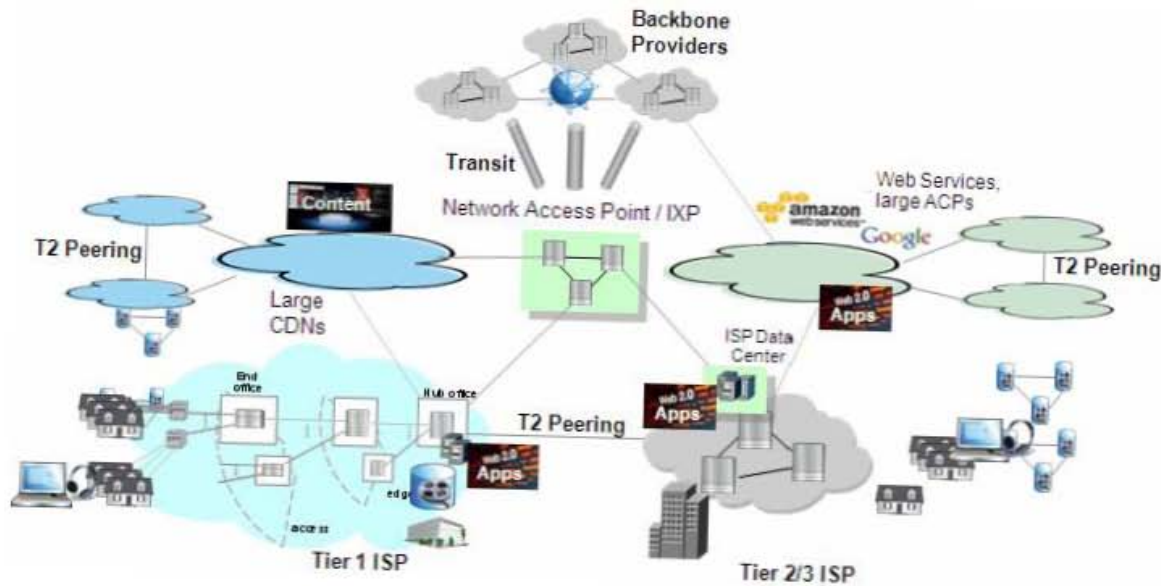


Figure 1. Schematic depiction of the ‘Internet’ as a set of interconnected managed networks

Internet Access services also traverse the operator network over the Access and Aggregation portions of the network, up to the IP edge, but at that point, they are connected to the Internet via an Internet ‘public peering point’ or ‘internet exchange’, as shown in Figure 1. Typically, quality of service or bandwidth guarantees are provided for the operator network Managed Services, however no guarantees are provided for generic Internet Access; the Internet Access service simply uses (a portion of) the bandwidth in the operator network, and competes with all other services running over the Internet.

The quality of service (QoS) distinction between Managed Services and Internet Access is, however, beginning to change as consumers begin to rely on services such as streaming video, communications and gaming services delivered as web or Internet services. For these services the application of quality of service is also desirable in order to receive the desired quality of experience (QoE).

In the following discussion, I therefore identify two types of Managed Services that I see as pertinent to the discussion of the end user Internet experience: (1) traditional ‘operator-selected’ Managed Services and (2) emerging ‘user-requested’ Managed Services. These two types of services are largely created using the same tools and features in the network, and may both use content delivery networks or connections to content providers with guaranteed connectivity (‘private peering’ connections) to allow the desired QoE to be delivered at all times.

Different Types/Categories of Managed Services

Broadly there are two categories of Managed Services:

- **Category 1:** those created by the operator for servicing an established need on the part of the end consumer, enterprise or PEG entity. These services have been traditionally been provided by service providers, e.g., IP television (IPTV), voice services and VPNs.
- **Category 2:** those requested directly by the end consumer, as a result of an enhanced QoS need for a specific service. This is an emerging category of service that is becoming increasingly important as end consumers seek to have a wider array of services, e.g. streaming, communications and gaming, made available with the same QoE as traditional Managed Services. It is important to note that for consumer-requested Managed Services, the request can be direct to the network operator, or can be indirect via a third party retail service provider (RSP) or application/content provider (ACP), acting on behalf of the end consumer.

It is important to understand that providing a higher bandwidth Best Effort service, e.g. in the form of a higher service tier to which the end user could subscribe, will not achieve the same desired quality of service. This will only typically change the peak bandwidth (the so-called Peak Information Rate or PIR) allocated to the subscriber over the Access provider's / ISP's network, but the services themselves will still compete with all other Best Effort services over the larger Internet, with no enhanced service treatment or differentiation. Thus, there is no real service guarantee, only a higher (but undefined) probability that the user will see higher average packet throughputs (but not necessarily result in lower delay or jitter, or any guaranteed throughput).

The majority of Category 1 services are provided entirely over managed IP networks; however, the emerging Category 2 services represent a far more diverse array of services since they are consumer-selected. As a result, these Category 2 services will also include services that are currently delivered over the Internet to the edge of the operator network, at which point the service will likely take advantage of the typical managed IP service delivery architecture, as well as dedicated peering connections or content delivery networks, as described above. For example, web content delivery services and "cloud computing" services are delivered over the Internet today, although increasingly these services are delivered via dedicated peering connections into operator networks to minimize the impact of the 'Best Effort' Internet on the service delivery.

An example of an Internet video streaming Managed Service is shown in Figure 2, whereby the content (file) is delivered to a Content Delivery Network (CDN) using a Best Effort service, but then is streamed to the end user from that cache over a peering connection into the managed IP network, with guaranteed Quality of Service.

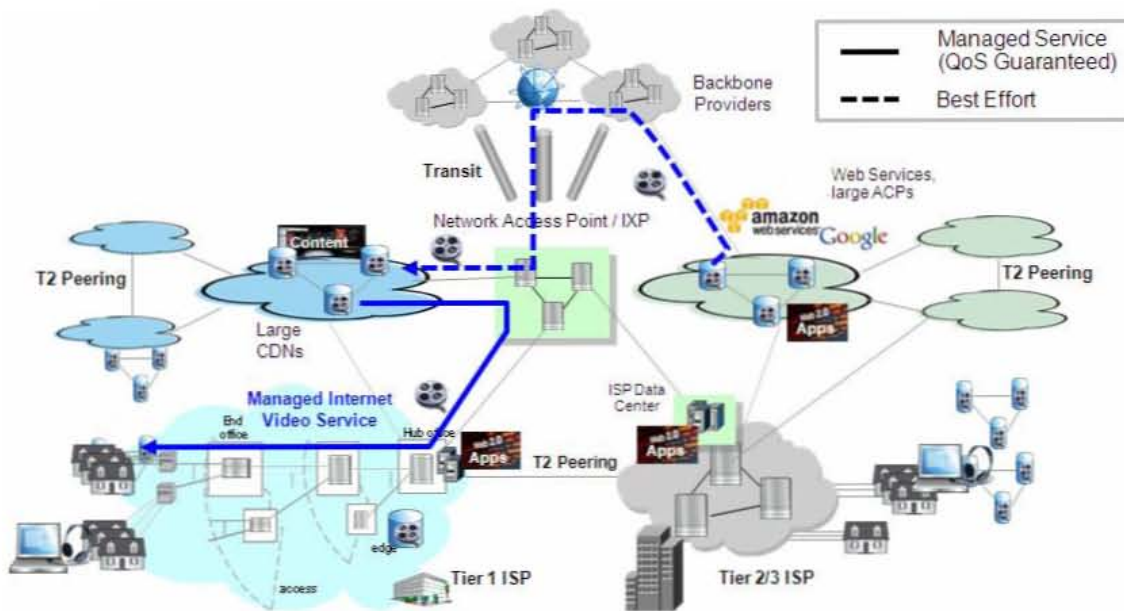


Figure 2. Depiction of the unmanaged and Managed Service domains

4 Managed Services Operation

Traffic Flow Control in Managed Service Offerings

There are many ways in which a Managed Service can provide control of the IP traffic flow, but typically this is done by controlling a number of parameters associated with the flow:

- The average guaranteed bandwidth (the so-called ‘Committed Information Rate’ or CIR)
- The average peak bandwidth (the ‘Peak Information Rate’ or PIR)
- The instantaneous maximum rate (the so-called ‘Maximum Burst Size’)

Packets are typically classified at the point of ingress to the Managed Service network, policed according to these parameters and then marked for forwarding in a specific forwarding class. This class is mapped to a specific queue type in the network element; the packets in the queue are then ‘de-queued’ according to well-defined algorithms that minimize the packet loss and delay associated with the different service flows.

This prioritized forwarding scheme always occurs independently of the load on the network, so the service delivery is always guaranteed. However, in the absence of excess network load (congestion), similar performance will also be seen in the “Best Effort” service class as, in this case, the queues in all forwarding classes (Managed and Best Effort) are emptied as fast as they are filled, with no concomitant packet loss experienced. Note that the delay (and therefore jitter) associated with the Best Effort service will likely be inferior to that achieved by the enhanced service class as there are

typically more services competing in the Best Effort class, using less well-controlled protocols (e.g. UDP), resulting in less predictable performance.

But, in effect, the goal of a Managed Service offering is therefore to guarantee the level of service at all times, whereas the Best Effort paradigm may deliver comparable performance but not at all times, i.e., not under all traffic load conditions or not for all applications simultaneously.

Operation of Managed Services

Traditionally, aggregate service management (e.g. CIR and PIR for the overall service usage across all users) occurs at ingress into the operator network with fine-grained (per-subscriber, per-service) policing occurring in the Broadband Remote Access Server (BRAS) or Broadband Network Gateway (BNG) element. However, increasingly, with the requirement for end-to-end QoS down to the end device and the increasing need for upstream bandwidth (from the consumer or enterprise), policing in the end consumer equipment and access network equipment is also required to ensure the requisite QoS.

For some services which can consume large amounts of bandwidth in a service tier or class, additional ‘services or network management’ may be required. For example, a fairness scheme might need to be applied to prevent one user’s traffic from utilizing the majority of a class shared with other users. This will apply whether this is a Managed Service class, or a Best Effort Internet service class.

Managed Services Are Responsive

The technical solutions outlined above respond to the two different categories of Managed Services and the interplay with Best Effort services, via interaction between the application and content domain with the network domain. This is typically achieved by a policy and admission control interface for dynamic services modification (in real time), or via management systems (for non-real time changes). In addition, other methodologies are possible, for example, using network and traffic optimization processes in the network as well as the movement of content and applications to the optimum location (e.g. using a CDN) for delivery to the end consumer that can complement QoS control methodologies, by optimizing the bandwidth available on different parts of the network. In these ways, it is possible to react to the changing applications and content mix, and requirements of end consumers for service guarantees, both during a given time period (e.g., minute, hour, day, and week) and over time (months or years).

Managed Services Configuration on Different Platforms

Managed Services can and should be configured using standard policy and admission control interfaces, as well as interfaces between the network management systems and the network elements. Some of these interfaces are vendor-specific, but they are currently being standardized or opened to third parties via well-defined protocols and

Application Programming Interfaces (APIs), allowing access to all the necessary capabilities by any third party. Furthermore, standardization will lead to the ability to offer the same type of service management on wireless, cable, satellite and wireline networks, and independent of the specific equipment vendor.

Managed Services Ecosystem

Today, there is relatively little interaction between the various parties in the broadband ecosystem in order to provide Managed Services. Typically, the network operator requires that the vendor support certain features that allow Managed Services (see QoS features referenced above, for example) to be provided with the desired scalability. The content or application provider closely collaborates with the network operator to ensure that the service can make use of the underlying network capabilities to provide a Managed Service.

Going forward, this paradigm will change to become much more open to any application developer or content provider, of any size or in any location, as the 'close collaboration' model will be extended to allow the Managed Services capabilities to be exposed via simple APIs, rather than a complex set of proprietary interfaces.

5 Managed Services Drive Innovation and Investment

There is a perception that Managed Services detract from or discourage investment in the network infrastructure because they use capacity that could otherwise be used by Best Effort services. However, my team and I have developed economic models which show that in fact the opposite is true: Managed Services enhance investment in the network due to the additional revenue potential they provide, and in turn, this enhances the capacity available for Best Effort services. The increase in overall network capacity *and* Best Effort capacity that results from the offering of Managed Services is clearly shown in Figure 3, which presents key results from our economic modeling of this issue.¹

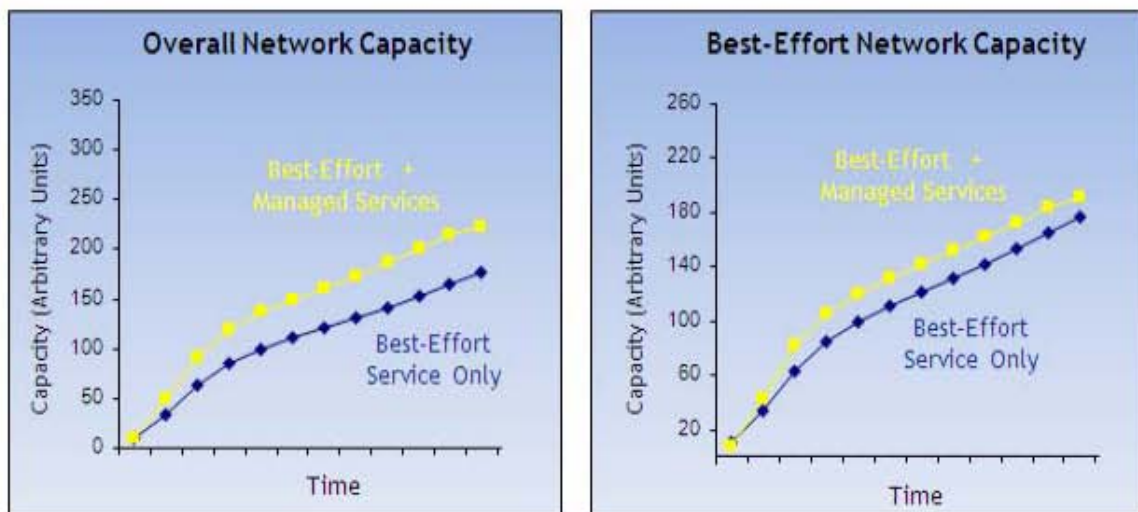


Figure 3. Results of an economic model of the interplay of Managed Services and Best Effort, showing the effect on overall network capacity (left) and the Best Effort network capacity (right)

Factors Shaping Investment in the Network

Network operators' network investment decisions are naturally tied to their expectations with respect to returns on that investment. American telecommunications providers have generally invested around 15-20% of their revenues over the last 5 years, based on publically disclosed financial information.ⁱⁱ Rules adopted in this proceeding, both generally and with specific regard to Managed Services, are likely to have a significant effect on investment decisions going forward. Because Managed Services offer unique value to individual and enterprise consumers and thus help ensure healthy revenues, the ability to offer such services increases expected returns on investment and promotes deployment. Moreover, because facilities used to provide Managed Services are also used to provide Best Effort Internet access service, the incentives created by the ability to offer Managed Services also advance investment in the "public Internet." Rules diminishing the provider's ability to design and offer services involving specialized, prioritized, and otherwise managed content would undercut these incentives, harming all broadband users.

As the FCC considers the proper treatment of managed services vis-à-vis Best Effort Internet service, it is also important to recognize that there is no single, predictable 'best' division between managed and unmanaged services or capacity, nor is there some constant, optimal investment figure. Rather, these factors are constantly in flux, driven by a large number of parameters. These include the evolving demand for Managed Service relative to Best Effort service (which relates to the relative quality of experience and the type of application, and the number of other applications competing for the capacity, and the demographics) and the changing cost of capacity (which relates to the network type, topography, and cost of equipment and labor). Because these factors are changing over time, the proper balance between managed and unmanaged service and the appropriate level of capital investment will change as well.

Regulatory Risks

We are entering a period of tremendous change in the definition of Managed Services, and in the associated business models – as noted, there is an emerging category of user-defined Managed Services offerings, and there is a manifold rate of increase in the diversity of services and applications available (and the change in the popularity of these applications in record time). Given this rapid rate of change, there is a very real risk that any attempts to explicitly and narrowly define what a 'Managed Service' is or to limit the number or variety of such services that are permitted, will seriously miss the mark and stifle innovation in the telecommunications industry. Moreover, given the European Union ruling that will support a diverse array of service management techniques,ⁱⁱⁱ the impact will be disproportionately felt in the U.S., which in turn, will lead to increased investment and revenue generation outside the U.S. and the concomitant exodus of top innovative talent.

The FCC should ensure that it defines Managed Services as broadly as possible to allow the continued growth of these important services, which will benefit both customers of Managed Services as well as users of the Best Effort Internet. Also, and as discussed in Section 7 below, the FCC should resist any effort to subject Managed Services to future network management rules since even the hint of a restrictive regulatory environment will limit the incentive to further develop Managed Services, which would ultimately impact the Best Effort Internet. Indeed, our economic modeling shows that the capacity growth can be artificially constrained if demand for Managed Services is fixed, for example, by a fixed regulatory specification of the size of the Managed Services bandwidth partition, and Best Effort services, and by extension, the overall social welfare, will suffer as a result.

6 The Evolution of the Managed Service Marketplace

New Types of Offerings

We can expect that the primary development in the Managed Service marketplace will be from the current operator-defined set of Managed Services, to also include a broader range of user-requested Managed Services, with an ever-changing set of applications and applications mash-ups required to be delivered as Managed Services.

New Management Tools

In addition, sophisticated dynamic, end-to-end policy management and admission control tools that apply across wireline, cable, satellite and wireless networks will become commonplace, and will complement new capabilities in management systems to optimize the network and allow per-user, per application/content QoS.

New Business Arrangements

The primary new business model will be that users will be able to freely select the applications that they require to receive Managed Services treatment from either the network operator or the ACP, who requests the service on behalf of the end user. The ACP will then enter into a business arrangement with the network operator to provide the requisite service for the requesting users. The net effect will be compelling new services delivered to the end users, and continual network investment on the part of the network operator.

7 Potential Implications of Broadband Industry Practice Regulations

Principle 1: Subject to reasonable network management, a provider of broadband Internet access service may not prevent any of its users from *sending or receiving the lawful content of the user's choice* over the Internet.

The application of Principle 1 to Managed Services (current or future) would negatively impact the value and utility of these services. Many Managed Services are offered outside of Best Effort high speed Internet because the content is extraordinarily valuable and/or sensitive. The service provider will often create the Managed Service specific to these concerns. For example, an IPTV service is offered on an IP network as a defined product – the service provider offers a variety of programs to the user with the QoS and security demanded by the content owner and end user. The ability to use the same IPTV service to allow access to content outside of the Managed Service could compromise the security and integrity of the service. Ultimately, whether third party content is permitted on a Managed Service should be at the discretion of the service provider, based on the network capabilities and user preferences, and as long as the decision is not anti-competitive.

Principle 2: Subject to reasonable network management, a provider of broadband Internet access service may not prevent any of its users from *running the lawful applications or using the lawful services of the user's choice*.

The application of Principle 2 to Managed Services (current or future) would negatively impact the value and utility of these services. A Managed Service is an end-to-end solution developed to offer the value demanded by the user. The service is application and/or service specific, and foreign applications and/or services, regardless of legality, could negatively impact the quality and dependability of the service. For example, a teleworking Managed Service enabling remote access may require applications to ensure the user can replicate specific work processes with unmitigated quality and security. The introduction of foreign applications, regardless of legality, may disrupt the teleworking Managed Service, which could compromise the security of the employer's network.

Principle 3: Subject to reasonable network management, a provider of broadband Internet access service may not prevent any of its users from *connecting to and using on its network the user's choice of lawful devices that do not harm the network*.

The application of Principle 3 to Managed Services (current or future) would negatively impact the value and utility of these services. A Managed Service is an end-to-end solution developed to offer the value demanded by the user. Part of this solution may require specific end user devices that support the Managed Service, such as an e-health solution that provides remote monitoring of a patient's health. This Managed Service may include specific applications and end user equipment to ensure the e-health service provides the functionality required with the requisite accuracy, quality and prioritization. The introduction of a foreign device to the e-health Managed Service could adversely affect the ability of this service to provide the quality, reliability and security that the end user demands. Ultimately, whether foreign devices may or may not be used with any Managed Service should be at the discretion of the service provider.

Principle 4: Subject to reasonable network management, a provider of broadband Internet access service may not deprive any of its users of the user's entitlement to *competition among network providers, application providers, service providers, and content providers*.

The application of Principle 4 to Managed Services (current or future) would negatively impact the value and utility of these services. A Managed Service is a product offered by the service provider to the end user. This product will be the result of numerous applications and service providers working with the service provider on a defined set of specifications. Enabling end users to introduce foreign applications or services to the Managed Service could compromise the quality and value of the service. For example, a smart grid Managed Service is offered to end users in order to better manage their electrical utility usage. The smart grid Managed Service is a finished product comprised of input from many different application and service providers. Entitling an end user with the ability to substitute any one of these applications or services vendors may compromise the integrity of the service.

Principle 5: Subject to reasonable network management, a provider of broadband Internet access service must *treat lawful content, applications, and services in a nondiscriminatory manner*.

The application of Principle 5 to Managed Services (current or future) would negatively impact the value and utility of these services. A Managed Service, by definition, offers the end user content, applications and services in a manner that has a higher level of QoS when compared to Best Effort broadband Internet access service. In order for this distinct QoS to be achieved, content, applications and services will have to be differentiated and treated distinctly. For example, the offering of a voice or streaming video service, which are sensitive to delays and packet loss, requires the service provider to prioritize this traffic over other traffic that is not delay or loss -sensitive. If such differentiation was prohibited, the managed voice or video service would fail to provide an adequate level of quality and the end user would not subscribe to such a service.

Principle 6: Subject to reasonable network management, a provider of broadband Internet access service must *disclose such information concerning network management and other practices* as is reasonably required for users and content, application, and service providers to enjoy the protections specified in this part.

The application of a broad transparency principle to Managed Services (current or future) would negatively impact the value and utility of these services. While general consumer disclosure typically is appropriate, in other cases, information concerning network management and other network practices should not be disclosed, particularly with respect to Managed Services. For example, the means by which a service provider ensures QoS and security for a content delivery or e-

health service should not be publicly available. These services are offered as complete products to end users and the network management tools employed by the service provider to ensure QoS and security may be proprietary or confidential. Moreover, the mandatory disclosure of such information could compromise the QoS and enhanced security required to enable a content delivery or e-health service. Ultimately, whether such network management information is disclosed should be at the discretion of the service provider.

ⁱ Bell Labs Mathematical Sciences Research Center, in preparation (2010).

ⁱⁱ *See, e.g.*, <http://www.google.com/finance?q=NYSE:T&fstype=ii>, and <http://www.google.com/finance?q=NYSE:VZ&fstype=ii>

ⁱⁱⁱ Council Directive 2009/136, 2009 O.J. (L337) 15 (“In order to meet quality of service requirements, operators may use procedures to measure and shape traffic on a network link so as to avoid filling the link to capacity or overfilling the link, which would result in network congestion and poor performance.”).